arXiv:1412.4463v2 [cs.DB] 27 Apr 2015

Defining relations on graphs: how hard is it in the presence of node partitions?

M Praveen and B Srivathsan
Chennai Mathematical Institute
Chennai, India

Abstract

Designing query languages for graph structured data is an active field of research. Evaluating a query on a graph results in a relation on the set of its nodes. In other words, a query is a mechanism for defining relations on a graph. Some relations may not be definable by any query in a given language. This leads to the following question: given a graph, a query language and a relation on the graph, does there exist a query in the language that defines the relation? This is called the definability problem. When the given query language is standard regular expressions, the definability problem is known to be PSPACE-complete.

The model of graphs can be extended by labeling nodes with values from an infinite domain. These labels induce a partition on the set of nodes: two nodes are equivalent if they are labeled by the same value. Query languages can also be extended to make use of this equivalence. Two such extensions are Regular Expressions with Memory (REM) and Regular Expressions with Equality (REE).

In this paper, we study the complexity of the definability problem in this extended model when the query language is either REM or REE. We show that the definability problem is EXPSPACE-complete when the query language is REM, and it is PSPACE-complete when the query language is REE. In addition, when the query language is a union of conjunctive queries based on REM or REE, we show coNP-completeness.

1 Introduction

Graph structures representing data have found many applications like semantic web [22, 15], social networks [23] and biological networks [18]. One model of graph structured data consists of a set of nodes labeled by values from some infinite domain and directed edges between the nodes labeled by letters from a finite alphabet. For example, a graph representing a social network may have a node for each member. There may be directed edges labeled *friend* between two nodes if the corresponding members are friends in the network. Nodes could be labeled by the name of the corresponding member's favourite movie. These labels from the infinite domain partition the set of nodes of the graph. Two nodes are equivalent if they have the same label. An active field of research is designing languages for querying such graphs, using both the structure of the graph and the partition induced by labels from the infinite domain [20, 4].

We will use the term *data graphs* for the model where nodes carry labels from an infinite domain (a nomenclature from [20]). The labels themselves are called *data values*. One way of querying data graphs is to simply specify a language L of strings. Each string in L has data values in odd positions and a letter from the finite alphabet in even positions. Evaluating a query specified by such a language on a data graph returns the set of all pairs of nodes $\langle q_1, q_2 \rangle$ such that there is a path from q_1 to q_2 labeled by a string in the specified language. Register automata [16, 24, 21] are extensions of standard finite state automata for handling data values from infinite domains. Using register automata as the formalism to specify languages, Libkin and Vrgoč studied the complexity of evaluating queries on data graphs [20]. There, the main reason behind the choice of register automata over other formalisms is to obtain tractable complexity for the query evaluation problem. Aiming towards a practically usable query language, extensions of standard regular expressions were introduced in [20]. They are named regular expressions

with memory (REM) and (less expressive) regular expressions with equality (REE). REM are expressively equivalent to register automata [19]. The complexity of query containment for these have also been studied [17].

Here we study the complexity of the definability problem: given a data graph and a set of pairs of nodes, check if the set can be obtained as the evaluation of some query on the data graph. One of the motivations for this study is the extraction of schema mappings, which we illustrate by an example. Given a data graph representing a social network, suppose we want to create another graph where two nodes are in the *movieLink* relation if they represent people having the same favourite movie and who are linked by a series of friends. There is a correspondence between the two graphs; in general such correspondences are called schema mappings. This particular schema mapping is specified by saying that the relation *movieLink* is exactly the relation returned by evaluating the query $friend^*$ on the original graph, with the additional condition that the two nodes have the same data value (i.e., the same favourite movie). Given the original data graph and the relation *movieLink*, suppose we want to algorithmically build the specification of the schema mapping using some query language. Then we need to check if the query language is capable of defining the *movieLink* relation — this is the definability problem. Using example instances of source and target schemas for deriving appropriate source-to-target mappings have been explored in relational databases [11, 14, 10, 2]. Research on schema mappings for graph databases has started [7, 5], though data values and extraction from example graphs have not been considered till now to the best of our knowledge. Example instances have also been used to derive "wrapper" queries for extraction of relevant information from data sources [13].

Contributions We study the complexity of the definability problem in data graphs, using either REM or REE as query languages. We prove the following results.

1. The definability problem with REM as the query language is EXPSPACE-complete.
2. The definability problem for REM with k memory locations is in $\text{SPACE}(\mathcal{O}(n\delta^k))$, where n is the number of nodes and δ is the number of data values used in the data graph.
3. The definability problem for REE is PSPACE-complete.
4. The definability problem for union of conjunctive queries based on REM or REE is coNP-complete.

For the upper bounds, we have to overcome some challenges. In the presence of data values, standard language theoretic tools like complementation, determinization and decidability of language inclusion do not work. We have to understand how data values affect definability, so that we can appeal directly to the more fundamental idea of pumping lemma, which still works in the presence of data values. For the lower bounds, we identify how small data graphs can count exponentially large numbers using data values, which otherwise require exponentially large graphs.

Related work Apart from derivation of mappings [11, 14, 10, 2], studies have also been made of using data examples to illustrate the semantics of schema mappings [1]. In [8], the problem of deriving schema mappings from data examples is studied from the perspective of algorithmic learning theory.

In [3], the complexity of the definability problem for graph query languages is studied, but they do not consider data values. Their main result is that definability using regular expressions in PSPACE-complete. They also give upper and lower bounds for various fragments of conjunctive queries based on regular expressions. We do not study conjunctive queries or their fragments but instead give tight bounds for union of conjunctive queries, which also apply to the setting of [3] where there are no data values.

The problem of query containment for fragments and extensions of REM and REE have been studied in [17]. A query e_1 is contained in another query e_2 if the set defined by e_1 is a subset of the set defined by e_2 on all data graphs. It is shown in [17] that for some fragments of REM and REE, query containment is respectively EXPSPACE-complete and PSPACE-complete. These are similar to the bounds we get for the definability problem. However, the upper bounds in [17] apply only to the positive fragments of REM and REE, where tests for inequality of data values are not allowed (query containment in the general case is undecidable). There is no obvious way to use those techniques here, since we allow the full syntax for REM and REE. For the EXPSPACE lower bound, the authors of [17] use techniques similar to those used in [6] to prove EXPSPACE lower bound for checking the emptiness of

parameterized regular expressions, closely related to REM. The EXPSPACE lower bound in [6] is based on succinctly reducing the emptiness of intersection of several expressions to the emptiness of a single expression. Here, we need to use a different approach, since we deal with the definability problem and can not rely on intersections.

2 Preliminaries

We will recall the basic definitions. The model of graphs with node labels from an infinite domain are called data graphs in [20]. We will follow the same nomenclature here. We will also make use of many other notations from [20].

Let Σ be a finite alphabet and let $\mathcal{D}$ be a countably infinite set of data values. We write $[n]$ for the set $\{0, 1, \ldots, n\}$.

Definition 1 (Data graph)**.** *A* data graph *over Σ and $\mathcal{D}$ is a triple $G = (V, E, \rho)$ where:*

- *V is a finite set of nodes,*
- *$E \subseteq V \times \Sigma \times V$ is a set of edges with labels in Σ,*
- *$\rho : V \to \mathcal{D}$ maps every vertex to a data value.*

Example 2. *Figure 1 gives an example of a data graph over $\Sigma = \{a\}$ and $\mathcal{D} = \mathbb{N}$, the set of natural numbers. However, a given graph would use only a finite set of data values. The role of data values will become clearer when we define query languages for such data graphs. We will use this graph as a running example throughout this section.*

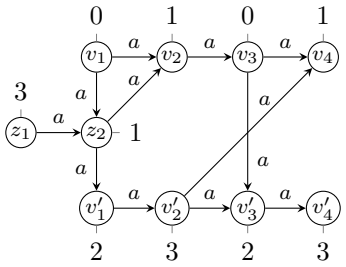

Figure 1: Example of a data graph over a unary alphabet $\Sigma = \{a\}$ and using data values $\{0, 1, 2, 3\}$.

A *path* in G is a sequence $\xi = v_1 a_1 v_2 a_2 \ldots v_{m-1} a_{m-1} v_m$ of nodes in V alternating with letters in Σ such that (v_i, a_i, v_{i+1}) is in E for all $i < m$. The *data path* w_ξ corresponding to a path ξ is the sequence $\rho(v_1)a_1\rho(v_2)a_2 \ldots \rho(v_{m-1})a_{m-1}\rho(v_m)$ obtained by replacing every node in ξ by its associated data value. We say that a data path w *connects* node u to v if there is a path $\xi = ua_1u_1 \ldots a_{m-1}v$ in G such that $w_\xi = w$. We write $u \xrightarrow{w} v$ in this case.

In general, a data path is a sequence $d_0a_0d_1a_1 \ldots a_{m-1}d_m$ of data values in $\mathcal{D}$ alternating with letters in Σ, starting and ending with data values. The set of all data paths over Σ and $\mathcal{D}$ is denoted by $\Sigma[\mathcal{D}]^*$. A *data language* $L \subseteq \Sigma[\mathcal{D}]^*$ is a set of data paths. Given two data paths $w_1 = d_0a_0d_1 \ldots a_{m-1}d_m$ and $w_2 = d'_0b_0d'_1 \ldots b_{l-1}d'_l$ where the last data value of w_1 coincides with the first data value of w_2 ($d_m = d'_0$), the *concatenation* $w_1 \cdot w_2$ is the data path $d_0a_1d_1 \ldots a_{m-1}d_mb_0d'_1 \ldots b_{l-1}d'_l$. This naturally extends to concatenations of many data paths. We will often write w_1w_2 instead of $w_1 \cdot w_2$.

We will now define two formalisms to characterize data languages. These formalisms would then be used to define query languages for data graphs. Since $\mathcal{D}$ could be infinite, these formalisms cannot check for the exact data value. They can however check for equality of two data values. The first formalism is *regular expressions with memory*. They are extensions of standard regular expressions over the finite alphabet Σ, introduced in [20]. They

are equipped with *registers*, that can store data values along a data path. The stored data values can be used to impose conditions on the data values allowed in future positions of the data path.

Definition 3. *Given a set of registers $r_1, r_2, \ldots, r_k$, the set $\mathcal{C}_k$ of conditions is given by the following grammar:*

$$c := \top \mid r_i^{=} \mid r_i^{\neq} \mid c \vee c \mid c \wedge c \mid \neg c, \quad 1 \leq i \leq k$$

The satisfaction is defined with respect to a data value $d \in \mathcal{D}$ and a tuple $\tau = (d_1, \ldots, d_k) \in (\mathcal{D} \cup \bot)^k$ called an assignment*: $d, \tau \models \top$ always, $d, \tau \models r_i^{=}$ iff $d_i = d$ and $d, \tau \models r_i^{\neq}$ iff $d_i \neq d$. The $\bot$ symbol is used to denote an empty register. It satisfies $\bot \neq d$ for every data value $d \in \mathcal{D}$. Satisfaction for the logical operators is as usual.*

Definition 4 (Regular expressions with memory)**.** *Let Σ be a finite alphabet and $r_1, \ldots, r_k$ a set of registers. Then,* regular expressions with memory (REM) *are defined by the following grammar:*

$$e := \varepsilon \mid a \mid e + e \mid e \cdot e \mid e^{+} \mid e[c] \mid \downarrow \overline{r}.e$$

where $a \in \Sigma$, $c \in \mathcal{C}_k$ and $\overline{r}$ is a tuple of registers.

We will use k-REM to denote the set of regular expressions with memory that use at most k registers. Let σ be an assignment. We will denote by $\sigma[\overline{r} \to d]$ the assignment obtained from σ by assigning d to the registers in $\overline{r}$. The semantics of k-REMs are as follows, reproduced from [20].

Definition 5 (Semantics of k-REMs)**.** *Suppose e is a k-REM, w is a data path in $\Sigma[\mathcal{D}]^*$ and $\sigma, \sigma' \in (\mathcal{D} \cup \bot)^k$ are assignments of the k registers used in e. The relation $(e, w, \sigma) \vdash \sigma'$ is defined by induction on the structure of e:*

$(\varepsilon, w, \sigma) \vdash \sigma'$	*if $w = d$ for some $d \in \mathcal{D}$ and $\sigma = \sigma'$*
$(a, w, \sigma) \vdash \sigma'$	*if $w = d_1 a d_2$ and $\sigma' = \sigma$*
$(e_1 + e_2, w, \sigma) \vdash \sigma'$	*if $(e_1, w, \sigma) \vdash \sigma'$ or $(e_2, w, \sigma) \vdash \sigma'$*
$(e_1 \cdot e_2, w, \sigma) \vdash \sigma'$	*if $w = w_1 \cdot w_2$ and $\exists\ \sigma_1 \in (\mathcal{D} \cup \bot)^k$ s. t. $(e_1, w_1, \sigma) \vdash \sigma_1$ and $(e_2, w_2, \sigma_1) \vdash \sigma'$*
$(e^{+}, w, \sigma) \vdash \sigma'$	*if $w = w_1 w_2 \ldots w_l$ and $\exists\ \sigma_0, \ldots, \sigma_l \in (\mathcal{D} \cup \bot)^k$ s.t. $(e, w_i, \sigma_i) \vdash \sigma_{i+1}$ for $i \in [l-1]$ and $\sigma_0 = \sigma$, $\sigma_l = \sigma'$*
$(e[c], w, \sigma) \vdash \sigma'$	*if $(e, w, \sigma) \vdash \sigma'$ and $\sigma', d \models c$ where d is the last data value in w*
$(\downarrow \overline{r}.e, w, \sigma) \vdash \sigma'$	*if $(e, w, \sigma[\overline{r} \to d]) \vdash \sigma'$ where d is the first value in w*

The *language* of a k-REM e is defined as follows:

$$\mathcal{L}(e) = \{\ w \in \Sigma[\mathcal{D}]^* \mid (e, w, \bot^k) \vdash \sigma \text{ for some } \sigma\}$$

where $\bot^k$ denotes the assignment that has $\bot$ in every register.

Example 6. *The REM $\downarrow r_1 \cdot a \cdot [r_1^{=}]$ uses one register. The language of this 1-REM consists of all data paths of the form dad where the first and last data values are the same. The 2-REM $\downarrow r_1 \cdot a \cdot \downarrow r_2 \cdot b \cdot a[r_1^{=}] \cdot b[r_2^{\neq}]$ contains data paths of the form $d_1 a d_2 b d_3 a d_4 b d_5$ where $d_1 = d_4$ and $d_2 \neq d_5$.*

The next formalism for characterizing data languages is another extension of standard regular expressions, called *regular expressions with equality*, again introduced in [20]. These are less powerful than REMs, since checking the equivalences between data values is restricted to a certain form.

Definition 7 (Regular expressions with equality)**.** *Let Σ be a finite alphabet and $\mathcal{D}$ a countably infinite set of data values. A* regular expression with equality (REE) *is constructed from the following grammar:*

$$e := \varepsilon \mid a \mid e + e \mid e \cdot e \mid e^+ \mid e_= \mid e_{\neq}$$

where a belongs to Σ. The language $\mathcal{L}(e)$ of an REE is defined as follows:

$$\begin{aligned}
\mathcal{L}(\varepsilon) &= \{\ d \mid d \in \mathcal{D}\} \\
\mathcal{L}(a) &= \{\ d_1 a d_2 \mid d_1, d_2 \in \mathcal{D}\} \\
\mathcal{L}(e_1 + e_2) &= \mathcal{L}(e_1) \cup \mathcal{L}(e_2) \\
\mathcal{L}(e_1 \cdot e_2) &= \mathcal{L}(e_1) \cdot \mathcal{L}(e_2) \\
\mathcal{L}(e^+) &= \{\ w_1 \cdots w_l \mid l \geq 1 \text{ and each } w_i \in \mathcal{L}(e)\ \} \\
\mathcal{L}(e_=) &= \{\ d_1 a_1 d_2 \ldots a_{m-1} d_m \in \mathcal{L}(e) \mid d_1 = d_m\} \\
\mathcal{L}(e_{\neq}) &= \{\ d_1 a_1 d_2 \ldots a_{m-1} d_m \in \mathcal{L}(e) \mid d_1 \neq d_m\}
\end{aligned}$$

Example 8. *The language of the REE $((a)_{\neq} \cdot (b)_{\neq})_{\neq}$ contains data paths $d_1 a d_2 b d_3$ such that $d_1 \neq d_2$, $d_2 \neq d_3$ and $d_1 \neq d_3$.*

We call a bijection $\pi : \mathcal{D} \to \mathcal{D}$ an automorphism on $\mathcal{D}$, since it preserves (in)equality.

Definition 9. *Let $\pi : \mathcal{D} \mapsto \mathcal{D}$ be an automorphism on $\mathcal{D}$. For a data path $w = d_0 a_0 d_1 a_1 \ldots d_m$ over $\Sigma[\mathcal{D}]^*$, we denote by $\pi(w)$ the data path $\pi(d_0) a_0 \pi(d_1) a_1 \ldots \pi(d_m)$ obtained by applying the automorphism π on the data values of w.*

An important property of REM and REE is that they cannot distinguish between automorphic data paths, just like register automata [16].

Fact 10 ([16, 19])**.** *For every REM or REE e, and for every data path $w \in \mathcal{L}(e)$ and automorphism $\pi : \mathcal{D} \to \mathcal{D}$, $\pi(w)$ is also in $\mathcal{L}(e)$.*

2.1 Query languages for data graphs

The above two formalisms can be used to define query languages for data graphs.

Definition 11 (Regular data path queries)**.** *An expression $Q = x \xrightarrow{e} y$ is a* regular data path query*, when e is either a standard regular expression, or an REM or an REE. Given a data graph G, the result of the query $Q(G)$ is the set of pairs of nodes $\langle u, v\rangle$ such that there exists a data path from u to v that belongs to $\mathcal{L}(e)$. The query is called* regular data path query with memory (RDPQ$_{mem}$) *or* regular data path query with equality (RDPQ$_=$) *depending on whether e is an REM or an REE. If e is a standard regular expression, the query is called a* regular path query (RPQ)*.*

A *relation* on the set of nodes in the graph is a set of tuples of same arity. We will say that a relation S on a data graph G is *defined* by a query Q if S equals $Q(G)$.

Example 12. *Evaluating the RPQ $Q_1 : x \xrightarrow{aaa} y$ on the data graph in Figure 1 results in the relation $S_1 = \{\ \langle v_1, v_4\rangle$, $\langle v_1, v'_3\rangle$, $\langle v_1, v_3\rangle$, $\langle v_1, v'_2\rangle$, $\langle v_2, v'_4\rangle$, $\langle z_1, v_3\rangle, \langle z_1, v'_2\rangle, \langle z_2, v_4\rangle$,*
$\langle z_2, v'_3\rangle, \langle v'_1, v'_4\rangle\ \}$. This is the set of all pairs of nodes connected by aaa. Neither $S_2 = \{\langle v_1, v_4\rangle, \langle v'_1, v'_4\rangle\}$ nor $S_3 = \{\langle v_1, v_3\rangle\}$ can be defined using RPQs. To see why, consider S_2. The only path connecting v'_1 to v'_4 is aaa.

But this path connects many other pairs apart from the ones in S_2. Hence to restrict to the pairs in S_2, we need to make use of data values. A similar argument will tell us that to define S_3, we need to consider data values.

The relation S_2 can be defined by the $\mathrm{RDPQ}_{\mathrm{mem}}$ $Q_2 : x \xrightarrow{e_2} y$, *where* $e_2 = \downarrow r_1 \cdot a \cdot \downarrow r_2 \cdot a[r_1^=] \cdot a[r_2^=]$. *The REM e_2 contains all data paths $d_1 a d_2 a d_3 a d_4$ such that $d_1 = d_3$ and $d_2 = d_4$. From Figure 1, it can be checked that the only data paths in the graph satisfying this expression are: $w_1 : 0a1a0a1$ and $w_2 : 2a3a2a3$, and they connect $\langle v_1, v_4 \rangle$ and $\langle v_1', v_4' \rangle$ respectively. Hence $Q_2(G) = S_2$, thus defining the relation S_2. Note that the two words w_1 and w_2 are automorphic images and hence cannot be distinguished by REMs (c.f. Fact 10).*

The expression e_2 is a 2-REM (uses 2 registers r_1 and r_2). Let us see why S_2 cannot be defined using a 1-REM. Suppose e is a 1-REM used in a query defining S_2. As the only data path connecting v_1' to v_4' is $2a3a2a3$, $\mathcal{L}(e)$ should contain the data path $2a3a2a3$. Moreover, the data paths $0a1a0a2$ and $1a2a3a2$ should not be in $\mathcal{L}(e)$, since $v_1 \xrightarrow{0a1a0a2} v_3'$ and $z_2 \xrightarrow{1a2a3a2} v_3'$. Since the prefix of $2a3a2a3$ (to be included) and $0a1a0a2$ (to be excluded) up to the first three data values are automorphic, the only way to add $2a3a2a3$ to $\mathcal{L}(e)$ and eliminate $0a1a0a2$ from $\mathcal{L}(e)$ is to check in e that the second and fourth data values are equal. This will still not eliminate $1a2a3a2$. To eliminate $1a2a3a2$, one has to add the condition that the first and third data values are equal. So we need to compare the first data value to the third, and second data value to the fourth. This kind of an "interleaved" check needs 2 registers as in the REM e_2 above. For the same reason, S_2 cannot be defined using $\mathrm{RDPQ}_{=}$.

The relation S_3 can be defined using the $\mathrm{RDPQ}_{=}$ $Q_3 : x \xrightarrow{e_3} y$ *with $e_3 = (a \cdot (a)_= \cdot a)_=$ In the data graph of Figure 1, the only data path satisfying e_3 is $w_5 : 0a1a1a0$. that connects $\langle v_1, v_3 \rangle$. Both the checks in e_3: first data value equals fourth data value, and second equals third, are required to eliminate the following words: $w_6 : 3a1a1a0$ and $w_7 : 1a2a3a1$, that connect $\langle z_1, v_3 \rangle$ and $\langle z_2, v_4 \rangle$ which are not in S_3. Hence for similar reasons as mentioned in the above paragraph, S_3 cannot be defined by an* $\mathrm{RDPQ}_{\mathrm{mem}}$ *that uses a 1-REM. A 2-REM would work though.*

We will also study a standard extension of query languages: union of conjunctive queries.

Definition 13 (Conjunctive data path queries)**.** *A* conjunctive regular data path query (CRDPQ) *is an expression of the form*

$$Ans(\overline{z}) := \bigwedge_{1 \le i \le m} x_i \xrightarrow{e_i} y_i, \tag{1}$$

where $m \geq 0$, x_i, y_i are variables and $\overline{z}$ is a tuple of variables among $\overline{x}$ and $\overline{y}$ and either every e_i is an REM, or every e_i is an REE. The semantics of a CRDPQ Q of the form (1) over a data graph $G = (V, E, \rho)$ is defined as follows. Given a valuation $\mu : \bigcup_{1 \le i \le m} \{x_i, y_i\} \to V$, we write $(G, \mu) \models Q$ if $\langle \mu(x_i), \mu(y_i) \rangle$ is in the answer of $x_i \xrightarrow{e_i} y_i$ on G, for each $i = 1, \ldots, m$. Then $Q(G)$ is the set of all tuples $\mu(\overline{z})$ such that $(G, \mu) \models Q$. The number of variables in $\overline{z}$ is the arity *of Q. A* union of conjunctive regular data path queries *(UCRDPQ) is a finite set $Q = \{Q_1, \ldots, Q_k\}$ of CRDPQs $Q_1, \ldots, Q_k$, which are all of the same arity. For a data graph G, $Q(G)$ is the set $Q_1(G) \cup \cdots \cup Q_k(G)$.*

Example 14. *We will work on the same graph from Figure 1. Consider the following CRDPQ Q_4: $Ans(x_1, y_1) := x_1 \xrightarrow{a} y_1 \ \wedge \ x_1 \xrightarrow{a} y_2 \ \wedge \ y_2 \xrightarrow{a} y_1$ The only valuation μ satisfying the above conditions is: $\mu(x_1) = v_1$, $\mu(y_1) = v_2$ and $\mu(y_2) = z_2$. The result $Q_4(G)$ would hence be the relation $\{ \langle v_1, v_2 \rangle \}$. Note that this relation cannot be defined using* $\mathrm{RDPQ}_{\mathrm{mem}}$ *or* $\mathrm{RDPQ}_{=}$. *The only data paths connecting $\langle v_1, v_2 \rangle$ are $0a1$ and $0a1a1$. The former data path cannot be used to distinguish $\langle v_1, v_2 \rangle$ as it connects $\langle v_3, v_4 \rangle$ as well and the latter one cannot be used since an automorphic data path $3a1a1$ connects $\langle z_1, v_2 \rangle$. From Fact 10, we know that REMs and REEs cannot differentiate between automorphic data paths.*

Consider another query Q_5: $Ans(x_1, y_1, x_2) := x_1 \xrightarrow{(a)_\neq} y_1 \ \wedge \ x_2 \xrightarrow{(a)_\neq} y_1$. The above query uses REEs in its individual regular data path queries. The result $Q_5(G)$ would be: $\{ \langle v_1, z_2, z_1 \rangle, \langle v_3, v_4, v_2' \rangle, \langle v_3, v_3', v_2' \rangle \}$ The query singles out the "pattern" of x_1 and x_2 converging into y_1, where the label of y_1 is different from those of x_1 and x_2.

2.2 Definability problems

From the examples, we can infer that RDPQ_{mem} and $\text{RDPQ}_{=}$ can define more relations than RPQ. In addition, RDPQ_{mem} can define more relations than $\text{RDPQ}_{=}$. CRDPQs can define even more than RDPQ_{mem}. Restricting to RDPQ_{mem}, using k registers we can define relations that are not possible with $k-1$ registers. It is also not difficult to construct examples of graphs and relations that are not definable using any of the query languages that we have seen. This motivates us to look at the following definability problems. The input is a data graph G and a relation S on the set of nodes in G.

RDPQ_{mem}-*definability:* Does there exist an RDPQ_{mem} Q s.t. $Q(G) = S$?

k-RDPQ_{mem}-*definability:* Does there exist an RDPQ_{mem} Q which uses at most k registers s.t. $Q(G) = S$?

$\text{RDPQ}_{=}$-*definability:* Does there exist an $\text{RDPQ}_{=}$ Q s.t. $Q(G) = S$?

UCRDPQ-definability: Does there exist a UCRDPQ Q s.t. $Q(G) = S$?

In the subsequent sections, we study the complexity of the above problems. For the last problem, we do not make a distinction between UCRDPQs using REM or REE, as we will see that the complexity stays the same in both cases.

Speciality of the equivalence relation As stated before, the data values induce an equivalence relation on the set of nodes, where two nodes are equivalent when they have the same label. Each letter a from the finite alphabet Σ also induces a binary relation on the set of nodes: the pair $\langle u, v\rangle$ is in this relation when there is an edge labeled a from u to v. Given that data values also induce a binary relation, why is it that we can not solve the definability problem by simply treating the equivalence relation as an extra letter in the finite alphabet and using the techniques developed for RPQs? The reason is that query languages give a special privilege to the equivalence relation: it can be used to relate positions that are far apart in a data path, while the binary relation induced by a letter in the finite alphabet can only relate successive positions. Hence, as seen in Example 12, some relations that can be defined by RDPQ_{mem} can not be defined by RPQ, even if we add the equivalence relation as an extra letter in the finite alphabet. However, a more sophisticated extension of the graph will allow us to use this idea, as explained in the beginning of the next section.

3 Queries using Regular Expressions with Memory

In this section we study the RDPQ_{mem}-definability and the k-RDPQ_{mem}-definability problems. Fix a data graph G and a binary relation S on the vertices in G. We denote the set of data values in G by $\mathcal{D}_G$. The goal is to decide if S is RDPQ_{mem}-definable. We start with some basic observations about the strengths and weaknesses of REMs.

If w is a data path and $\pi : \mathcal{D}_G \to \mathcal{D}_G$ is an automorphism on $\mathcal{D}_G$, we have seen in Fact 10 that no REM can distinguish between w and $\pi(w)$. On the other hand, if two data paths are not automorphic, then they can be distinguished by an REM. We denote by $[w]$ the set of all data paths automorphic to w.

Lemma 15. *For every data path w, there is an REM $e_{[w]}$ such that $\mathcal{L}(e_{[w]}) = [w]$.*

Proof. Suppose $d_1, \ldots, d_k$ are the distinct data values occurring in w. The required REM $e_{[w]}$ uses k registers $r_1, \ldots, r_k$. Essentially, at the first position where the data value d_i appears, $e_{[w]}$ stores d_i in the register r_i. In every subsequent position where d_i appears, it is compared against the value stored in r_i. Formally, $e_{[w]}$ is defined

as follows by induction on length of w: $e_{[d_i]} = \downarrow r_i$.

$$e_{[wad_i]} = \begin{cases} e_{[w]} \cdot a[r_i^=] & \text{if } d_i \text{ occurs in } w \\ e_{[w]} \cdot a \cdot \downarrow r_i.\varepsilon & \text{otherwise} \end{cases}$$

For every data value d_i occurring in w, $e_{[w]}$ notes all the positions having the data value d_i. Using this, it is routine to prove that any data path w' is in $\mathcal{L}(e_{[w]})$ iff w' is automorphic to w. □

Combining Fact 10 and Lemma 15, we infer that two data paths can be distinguished by an REM iff they are not automorphic. This suggests the following procedure for checking RDPQ_{mem}-definability. For simplicity, suppose that we want to define the singleton set $\{\langle u, v\rangle\}$. Suppose there is a data path w connecting u to v. The expression $e_{[w]}$ will not define the set $\{\langle u, v\rangle\}$ iff there is an automorphism π such that $\pi(w)$ connects u' to v' for some $\langle u', v'\rangle \neq \langle u, v\rangle$. The automorphism π is obstructing $e_{[w]}$ from defining $\{\langle u, v\rangle\}$, but this obstruction is not explicit in the data graph G. It *is* explicit in $G_{\pi^{-1}}$ (obtained from G after replacing every data value d by $\pi^{-1}(d)$), since w connects u' to v' in $G_{\pi^{-1}}$. All such obstructions will be explicit in G_{aut}, the disjoint union of G_π for all automorphisms π. A little more work will allow us to drop the special treatment given to data values and treat them as usual letters from a finite alphabet in G_{aut}. The RDPQ_{mem}-definability problem on G can be reduced to the RPQ-definability problem on G_{aut}. The PSPACE-completeness of RPQ-definability [3] will then give an EXPSPACE upper bound for RDPQ_{mem}-definability. This approach however does not throw light on the role of registers in definability, nor does it give precise bounds in the case where the number of registers is fixed.

In the next sub-section, we make some observations on $k - \text{RDPQ}_{\text{mem}}$-definability, which are counterparts of the above observations on RDPQ_{mem}-definability.

3.1 RDPQ_{mem}-definability with bounded number of registers

If w is a data path with k distinct data values, we saw in Lemma 15 that there is a REM $e_{[w]}$ whose language is exactly $[w]$. The number of registers used in $e_{[w]}$ is k. If we restrict the number of registers to less than k, then there may not be an expression whose language is exactly $[w]$. Still, the expression $e_{[w]}$ (which uses k registers) has a simple syntactic form, which we would like to capture and use in scenarios where there are fewer registers.

Definition 16 (Basic REM)**.** *A* basic k-REM *is a k-REM of the form $\downarrow \overline{r}_1.a_1[c_1] \cdot \downarrow \overline{r}_2.a_2[c_2] \cdots \downarrow \overline{r}_m.a_m[c_m]$, where $a_i \in \Sigma$, $c_i \in \mathcal{C}_k$ and $\overline{r}_i$ are tuples from $r_1, \ldots, r_k$.*

Basic k-REMs can also be thought of as those built without using the rules $e := e^+$ and $e := e+e$. We considered defining a singleton set $\{\langle u, v\rangle\}$ for simplicity. We would like to retain the simplicity but handle arbitrary sets, which is the purpose of the following definition.

Definition 17. *Suppose G is a data graph, S is a binary relation on the set of nodes of G and $\langle u, v\rangle \in S$. A k-REM witness for $\langle u, v\rangle$ in S is a basic k-REM e satisfying the following conditions.*

1. *(Connecting path) $u \xrightarrow{w} v$ for some $w \in \mathcal{L}(e)$.*
2. *(No extraneous pairs) If any data path in $\mathcal{L}(e)$ connects some u' to some v', then $\langle u', v'\rangle \in S$.*

If an arbitrary k-REM e defines S and $\langle u, v\rangle \in S$, then there is a data path $w \in \mathcal{L}(e)$ connecting u to v. If e is of the form $e_1 e_2^+ e_3$, there is an m such that $\mathcal{L}(e_1 e_2^m e_3)$ contains w. Continuing this process of removing iterations in e, while still retaining w in the language will result in a k-REM witness for $\langle u, v\rangle$ in S.

Lemma 18. *If S is definable, then it is definable by a union of $|S|$ k-REM witnesses.*

Proof. Suppose a k-REM e defines S and $\langle u, v\rangle \in S$. We will show that there exists a k-REM witness for $\langle u, v\rangle$ in S, which will prove the lemma.

Since e defines S, there is a data path $w \in \mathcal{L}(e)$ connecting u to v. Without loss of generality, we can assume that e is of the form $e_1 + e_2 + \cdots + e_m$ for some $m \geq 1$ such that each e_i is union-free, that is each e_i is constructed using the grammar for REMs without the $e := e + e$ rule. The data path w belongs to $\mathcal{L}(e_i)$ for some $i \in \{1, \ldots, m\}$. If e_i is a basic k-REM, then we are done. Otherwise, e_i is of the form $f_1 \cdot (f_2)^+ \cdot f_3$ where f_1, f_2 and f_3 are union free k-REMs. As $w \in \mathcal{L}(e_i)$, from the semantics of k-REMs, there exists a number α such that w satisfies the k-REM obtained by α iterations of f_2. More precisely, there exist data paths $x \in \mathcal{L}(f_1)$, $y_1, \ldots, y_\alpha \in \mathcal{L}(f_2)$ and $z \in \mathcal{L}(f_3)$ such that $w = x \cdot y_1 \cdots y_\alpha \cdot z$.

Let us write $(f_2)^\alpha$ for the k-REM obtained by concatenating f_2 α times. If $f_1(f_2)^\alpha f_3$ is basic, then we are done. Otherwise continue this process of "unfolding" to get a basic k-REM e'. Since w is in $\mathcal{L}(e')$, e' satisfies the first condition of Definition 17 (connecting path). Since $\mathcal{L}(e') \subseteq \mathcal{L}(e)$ and e defines S, e' satisfies the second condition of Definition 17 (no extraneous pairs). Hence, e' is a k-REM witness for $\langle u, v \rangle$ in S, which finishes the proof. □

Now suppose we are trying to define S using $k - \mathrm{RDPQ}_{\mathrm{mem}}$ and let $\langle u, v \rangle \in S$. Assume there is a data path w connecting u to v and there is a basic k-REM e such that $w \in \mathcal{L}(e)$. If e is not a k-REM witness for $\langle u, v \rangle$ in S, then there is a data path $w' \in \mathcal{L}(e)$ connecting u' to v' for some $\langle u', v' \rangle \neq \langle u, v \rangle$. The data path w' is obstructing e from being a witness and we need a structure where such obstructions are explicit. Since we are dealing with k-REMs, the structure would have to keep track of possible values stored in the k registers. The following definition and lemma are similar to the way the semantics of REM over a data graph is defined in [17].

Definition 19 (Assignment graph)**.** *Let k be a natural number. To a data graph $G = (V, E, \rho)$ over finite alphabet Σ and data values $\mathcal{D}_G$ we associate a transition system $\mathcal{T}_G = (Q_G, \to_G)$ called the* k-assignment graph. *Its set of states is $Q_G = V \times (\mathcal{D}_G \cup \bot)^k$. The transitions are of the form $\downarrow \overline{r}.a[c]$, where $\overline{r}$ is a (possibly empty) tuple of variables from $r_1, \ldots, r_k$, $a \in \Sigma$ and c is a condition in $\mathcal{C}_k$. There is a transition $(v, \sigma) \xrightarrow{\downarrow \overline{r}.a[c]}_G (v', \sigma')$ if $(v, a, v') \in E$, $\sigma' = \sigma[\overline{r} \to \rho(v)]$ and $\rho(v'), \sigma' \models c$.*

A sequence of the form $(v_0, \sigma_0) \xrightarrow{\downarrow \overline{r}_1.a_1[c_1]}_G (v_1, \sigma_1) \to_G \cdots \xrightarrow{\downarrow \overline{r}_m.a_m[c_m]}_G (v_m, \sigma_m)$ in $\mathcal{T}_G$ is called a run from (v_0, σ_0) to (v_m, σ_m). The sequence $\downarrow \overline{r}_1.a_1[c_1] \cdots \downarrow \overline{r}_m.a_m[c_m]$ is a basic k-REM. Hence, we can think of runs in $\mathcal{T}_G$ as being of the form $(u, \sigma) \xrightarrow{e}_G (v, \sigma')$, where e is the basic k-REM formed by the labels of the sequence of transitions connecting (u, σ) to (v, σ'). This observation leads to the following connection between runs in the assignment graph and data paths in G belonging to the languages of basic REMs.

Lemma 20. *Let e be a basic k-REM. Let $\sigma : \{r_1, \ldots, r_k\} \to \mathcal{D}_G \cup \{\bot\}$ and $\sigma' : \{r_1, \ldots, r_k\} \to \mathcal{D}_G \cup \{\bot\}$ be some assignments. The following are equivalent.*

1. *A data path w connects u to v in G and $(e, w, \sigma) \vdash \sigma'$.*
2. *There exists a run $(u, \sigma) \xrightarrow{e}_G (v, \sigma')$ in $\mathcal{T}_G$.*

Proof. By an induction on the number of blocks of the form $\downarrow \overline{r}.a[c]$ in e. Suppose $e = \downarrow \overline{r}.a[c]$. If there is a data path w as in the lemma, then $(u, a, v) \in E$ and by the semantics of REMs, $\sigma' = \sigma[\overline{r} \to \rho(u)]$ and $\rho(v), \sigma' \models c$. Hence $(u, \sigma) \xrightarrow{\downarrow \overline{r}.a[c]}_G (v, \sigma')$. Conversely, let us suppose $(u, \sigma) \xrightarrow{\downarrow \overline{r}.a[c]}_G (v, \sigma')$ in $\mathcal{T}_G$. By definition of $\mathcal{T}_G$, we have $(u, a, v) \in E$, $\sigma' = \sigma[\overline{r} \to \rho(u)]$ and $\rho(v), \sigma' \models c$. Hence, $\rho(u)a\rho(v)$ is a data path connecting u to v in G and $(\downarrow \overline{r}.a[c], \rho(u)a\rho(v), \sigma) \vdash \sigma'$.

For the induction step, suppose $e = \downarrow \overline{r}.a[c] \cdot e'$. If there is a data path w as in the lemma, then $w = \rho(u)a\rho(u_0) \cdot w'$ for some node u_0 and a data path w' connecting u_0 to v in G. In addition, by the semantics of REMs, there is some assignment σ_0 such that $(\downarrow \overline{r}.a[c], \rho(u)a\rho(u_0), \sigma) \vdash \sigma_0$ and $(e', w', \sigma_0) \vdash \sigma'$. Now we can use an argument similar to the one in the base case to infer that $(u, \sigma) \xrightarrow{\downarrow \overline{r}.a[c]}_G (u_0, \sigma_0)$ and use the induction hypothesis to infer that $(u_0, \sigma_0) \xrightarrow{e'}_G (v, \sigma')$. Hence, $(u, \sigma) \xrightarrow{e}_G (v, \sigma')$. Conversely, suppose $(u, \sigma) \xrightarrow{e}_G (v, \sigma')$ in $\mathcal{T}_G$. This run can be split as follows: $(u, \sigma) \xrightarrow{\downarrow \overline{r}.a[c]} (u_0, \sigma_0) \xrightarrow{e'}_G (v, \sigma')$ for some node u_0 and assignment σ_0. Then we can argue as in the base case to infer that the data path $\rho(u)a\rho(u_0)$ connects u to u_0 in G and $(\downarrow \overline{r}.a[c], \rho(u)a\rho(u_0), \sigma) \vdash \sigma_0$. We can use the induction hypothesis to infer that there is a data path w' connecting u_0 to v in G and $(e', w', \sigma_0) \models \sigma'$. Hence, the data path $\rho(u)a\rho(u_0) \cdot w'$ connects u to v in G and $(e, \rho(u)a\rho(u_0) \cdot w', \sigma) \vdash \sigma'$. □

Suppose we are trying to define a set S on the data graph G using $k - \mathrm{RDPQ}_{\mathrm{mem}}$. The above lemma allows us to think of k-REM witnesses in terms of runs in $\mathcal{T}_G$. A basic k-REM e is a k-REM witness for $\langle u, v\rangle$ in S iff it satisfies the following conditions.

1. $(u, \bot^k) \xrightarrow{e}_G (v, \sigma)$ for some assignment σ, to satisfy condition 1 of Definition 17 (connecting path).
2. If $(u', \bot^k) \xrightarrow{e}_G (v', \sigma)$ for some nodes u', v' and some assignment σ, then $\langle u', v'\rangle \in S$, to satisfy condition 2 of Definition 17 (no extraneous pairs).

Checking that a basic k-REM e is a witness thus reduces to checking that e connects a pair in $\mathcal{T}_G$ and does not connect certain other pairs. This observation allows us to use the pigeon hole principle to prove the existence of short witnesses.

Lemma 21. *Suppose G is a data graph with δ distinct data values, n nodes $v_1, \ldots, v_n$ and S is a binary relation on the set of nodes. If there is a k-REM witness for $\langle v_p, v_q\rangle$ in S, there is one of length $\mathcal{O}\left(2^{n^2\delta^k}\right)$.*

Proof. For sets of states $Q_1, \ldots, Q_n, Q'_1, \ldots, Q'_n \subseteq Q_G$, we write $\langle Q_1, \ldots, Q_n\rangle \xrightarrow{e}_G \langle Q'_1, \ldots, Q'_n\rangle$ if $Q'_i = \{(v', \sigma') \mid (v, \sigma) \xrightarrow{e}_G (v', \sigma')$ for some $(v, \sigma) \in Q_i\}$ for every $i = 1, \ldots, n$. Suppose e is a k-REM witness for $\langle v_p, v_q\rangle$ in S. Let $e = e_1 \cdot e_2 \cdots e_m$, where every e_i is of the form $\downarrow \overline{r}_i a_i.[c_i]$. Consider the sequence:

$$\langle \{(v_1, \bot^k)\}, \ldots, \{(v_n, \bot^k)\}\rangle \xrightarrow{e_1}_G \langle Q^1_1, \ldots, Q^1_n\rangle \xrightarrow{e_2}_G \cdots \xrightarrow{e_m}_G \langle Q^m_1, \ldots, Q^m_n\rangle \ . \tag{2}$$

The set Q^j_i is the set of all states reachable from $(v_i, \bot^k)$ along the path $e_1 \cdots e_j$ in $\mathcal{T}_G$. If there are $j < j'$ such that $\langle Q^j_1, \ldots, Q^j_n\rangle = \langle Q^{j'}_1, \ldots, Q^{j'}_n\rangle$, then removing the part of this sequence between j and j' will lead to the same final tuple $\langle Q^m_1, \ldots, Q^m_n\rangle$. We claim that after this removal, the resulting k-REM $e_1 \cdots e_j \cdot e_{j'+1} \cdots e_m$ is a k-REM witness for $\langle v_p, v_q\rangle$ in S. The reason is as follows: from Lemma 20, the following two conditions are equivalent to the original hypothesis that e is a k-REM witness for $\langle v_p, v_q\rangle$ in S.

1. For some assignment σ, $(v_q, \sigma) \in Q^m_p$.
2. For any $i = 1, \ldots, n$ and any $(v, \sigma) \in Q^m_i$, $(v_i, v) \in S$.

Hence, any basic k-REM that ends in the same n-tuple $\langle Q^m_1, \ \ldots \ , Q^m_n\rangle$ is also a k-REM witness for $\langle v_p, v_q\rangle$ in S. As long as there are duplicate tuples along the sequence (2), we can remove part of it to get a shorter witness. By pigeon hole principle, we conclude that there is a witness no longer than the total number of distinct tuples $\langle Q_1, \ldots, Q_n\rangle$.

There are at most $n(\delta + 1)^k$ states in $\mathcal{T}_G$. Hence, there are at most $2^{n^2(\delta+1)^k}$ tuples $\langle Q_1, \ldots, Q_n\rangle$. From the argument in the previous paragraph, we infer that if there is a k-REM witness for $\langle v_p, v_q\rangle$ in S, there is one of length at most $2^{n^2(\delta+1)^k}$. □

For graphs without data considered in [3], the solution to RPQ-definability looks at the graph as a finite automaton. This paves the way for using language theoretic tools, which are ultimately based on a pumping argument. In our case, we cannot view a data graph directly as a register automaton. Hence we need to construct the assignment graph on which we can apply the pumping argument.

Theorem 22. *The $k - \mathrm{RDPQ}_{\mathrm{mem}}$-definability problem is in $\mathrm{NSPACE}(\mathcal{O}(n^2\delta^k))$, where n is the number of nodes and δ is the number of distinct data values.*

Proof. Suppose we are trying to define the set S. From Lemma 18, it is enough to check that there are $|S|$ k-REM witness, one for each pair $\langle u, v\rangle$ in S. From Lemma 21, we infer that it is enough to check for witnesses of length at most $2^{n^2(\delta+1)^k}$. We will now give a non-deterministic algorithm to do this in space $\mathcal{O}(n^2\delta^k)$.

First we note that given a tuple $\langle Q_1, \ldots, Q_n \rangle$ of subsets of Q_G and a k-REM $e =\downarrow \overline{r}a.[c]$, we can compute in space polynomial in $(n\delta^k)$ the tuple $\langle Q'_1, \ldots, Q'_n \rangle$ such that $\langle Q_1, \ldots, Q_n \rangle \xrightarrow{e}_G \langle Q'_1, \ldots, Q'_n \rangle$. Suppose $v_1, \ldots, v_n$ are the nodes of G. Now we give a non-deterministic algorithm to check if there exists a k-REM witness for $\langle v_1, v_p \rangle$ in S of length at most $2^{n^2(\delta+1)^k}$. The algorithm maintains a counter initialized to 0 and a tuple of subsets of Q_G, initialized to $\langle \{(v_1, \perp^k)\}, \ldots, \{(v_n, \perp^k)\} \rangle$. The algorithm performs the following steps as long as the counter does not exceed $2^{n^2(\delta+1)^k}$.

1. Increment the counter.
2. Guess a k-REM $e =\downarrow \overline{r}a.[c]$.
3. Replace current tuple $\langle Q_1, \ldots, Q_n \rangle$ with $\langle Q'_1, \ldots, Q'_n \rangle$, where $\langle Q_1, \ldots, Q_n \rangle \xrightarrow{e} \langle Q'_1, \ldots, Q'_n \rangle$.
4. Check if $(v_p, \sigma) \in Q'_1$ for some assignment σ and that for every $i = 1, \ldots, n$ and for every $(v', \sigma) \in Q'_i$, the pair $\langle v_i, v' \rangle$ belongs to S. If yes, accept and terminate. If not, go back to step 1.

From the proof of Lemma 21, we conclude that some run of the above non-deterministic algorithm will accept if there is a k-REM witness for $\langle v_1, v_p \rangle$ in S. If there is no such witness, then clearly no run will accept. The algorithm needs space to store the counter, the tuple of subsets of Q_G and the space to compute the successor tuple. The counter can be implemented in space $\mathcal{O}(n^2\delta^k)$ using binary counting. One state of Q_G needs $(\log n \cdot k \cdot \log \delta)$ bits. There are at most $n(\delta+1)^k$ states in Q_G. Hence, the tuple of subsets and the space needed for intermediate computations can all be accommodated in space $\mathcal{O}(n^2\delta^k)$. □

3.2 $\mathrm{RDPQ}_{\mathrm{mem}}$-definability

We can now tackle $\mathrm{RDPQ}_{\mathrm{mem}}$-definability, where there is no bound on the number of registers.

Lemma 23. *Suppose G is a data graph with δ distinct data values. A relation S is $\mathrm{RDPQ}_{\mathrm{mem}}$-definable if and only if it is δ-$\mathrm{RDPQ}_{\mathrm{mem}}$-definable.*

Proof. The right to left implication is obvious. For the other direction, suppose S is $k-\mathrm{RDPQ}_{\mathrm{mem}}$-definable for some k. From Lemma 18, for every pair $\langle u, v \rangle \in S$, there is a k-REM witness e for $\langle u, v \rangle$ in S. Hence, there exists a data path $w \in \mathcal{L}(e)$ connecting u to v. From Lemma 15, $e_{[w]}$ is a δ-REM and from Fact 10, $\mathcal{L}(e_{[w]}) \subseteq \mathcal{L}(e)$. Hence, $e_{[w]}$ is a δ-REM witness for $\langle u, v \rangle$ in S. Such witnesses exist for every pair in S and hence, S is $\delta-\mathrm{RDPQ}_{\mathrm{mem}}$-definable. □

The next theorem follows from the previous two results.

Theorem 24. $\mathrm{RDPQ}_{\mathrm{mem}}$*-definability is in* Expspace*.*

Proof. From Lemma 23, it is equivalent to checking $\delta - \mathrm{RDPQ}_{\mathrm{mem}}$-definability, where δ is the number of distinct data values in the given data graph. From Theorem 22, $\delta - \mathrm{RDPQ}_{\mathrm{mem}}$-definability is in $\mathrm{NSpace}(\mathcal{O}(n^2\delta^\delta))$. From Savitch's theorem, we then get a deterministic exponential space algorithm. □

Next we give a matching lower bound.

Theorem 25. *The $\mathrm{RDPQ}_{\mathrm{mem}}$-definability problem in data graphs is* Expspace*-hard.*

Proof. We reduce the exponential width corridor tiling problem to the $\mathrm{RDPQ}_{\mathrm{mem}}$-definability problem. An instance of the tiling problem consists of a set T of tile types, a relation $C_h \subseteq T \times T$ of horizontally compatible tile types and a relation $C_v \subseteq T \times T$ of vertically compatible tile types, an initial tile type t_i, a final tile type t_f and a number n (in unary). The problem is to check if there exists a number R and a tiling $\tau : [R] \times [2^n - 1] \to T$ that is *legal* — $\tau(0,0) = t_i$, $\tau(R, 2^n - 1) = t_f$, $(\tau(i,j), \tau(i,j+1)) \in C_h$ and $(\tau(i,j), \tau(i+1,j)) \in C_v$ for all i, j. The intention here is that $\tau(i,j)$ is the tile type at the i^{th} row j^{th} column of a corridor with $R+1$ rows and 2^n columns. This problem

is known to be EXPSPACE-complete (e.g., see [25]). To be precise, we need to allow any exponential function in place of 2^n. Our proof works in that case also; we use 2^n to reduce notational clutter.

Let $\overline{T} = \{\overline{t} \mid t \in T\}$ be a disjoint copy of T. Given an instance of the tiling problem, we reduce it to the RDPQ_{mem}-definability problem in data graphs, where the finite alphabet is $T \cup \overline{T} \cup \{\$, \alpha\}$. A tiling τ is encoded by data paths in the language of the following REM:

$$
\begin{array}{llllll}
\$\cdot & \downarrow r_n \cdot & \alpha\cdot & \downarrow r_{n-1}\cdot & \alpha\cdots\alpha\cdot & \downarrow r_1\cdot \quad \tau(0,0) \\
 & [r_n^{=}]\cdot & \alpha & [r_{n-1}^{=}]\cdot & \alpha\cdots\alpha & [r_1^{\neq}]\cdot \quad \tau(0,1) \\
 & [r_n^{=}]\cdot & \alpha & \cdots\alpha & [r_2^{\neq}]\cdot\alpha & [r_1^{=}]\cdot \quad \tau(0,2) \\
 & & & & \vdots & \\
 & [r_n^{\neq}]\cdot & \alpha & [r_{n-1}^{\neq}]\cdot & \alpha\cdots\alpha & [r_1^{\neq}]\cdot \quad \overline{\tau(0,2^n-1)} \\
 & [r_n^{=}]\cdot & \alpha & [r_{n-1}^{=}]\cdot & \alpha\cdots\alpha & [r_1^{=}]\cdot \quad \tau(1,0) \\
 & & & & \vdots & \\
 & [r_n^{\neq}]\cdot & \alpha & [r_{n-1}^{\neq}]\cdot & \alpha\cdots\alpha & [r_1^{\neq}]\cdot \quad \overline{\tau(R,2^n-1)}\cdot \quad \$
\end{array}
\tag{3}
$$

The expression lists the tile types used in the tiling sequentially from left column to right column, bottom row to top row. The first n data values are stored in the registers $r_n, \ldots, r_1$. In later positions, $[r_k^{=}]$ (resp. $[r_k^{\neq}]$) indicates that the k^{th} bit is 0 (resp. 1). The n conditions preceding $\tau(i,j)$ in the expression denote the binary representation of j. Tile types in the last column are represented by letters in $\overline{T}$, so that we need not check them for horizontal compatibility with the next tile. The data graph is the disjoint union of two graphs $p_1 \xrightarrow{\$} \boxed{\text{illegal tilings}} \xrightarrow{\$} q_1$ and $p_2 \xrightarrow{\$} \boxed{\text{all tilings}} \xrightarrow{\$} q_2$ satisfying the following conditions.

1. Any data path starting and ending with the letter \$ may only connect p_1 to q_1 or p_2 to q_2.
2. Every tiling can be encoded by some data path connecting p_2 to q_2.
3. None of the data paths connecting p_1 to q_1 are encodings of legal tilings.
4. For every data path w connecting p_2 to q_2 that is not the encoding of a legal tiling, there exists a data path automorphic to w connecting p_1 to q_1.

We claim that there exists a legal tiling iff $\{\langle p_2, q_2\rangle\}$ is RDPQ_{mem}-definable. Indeed, suppose there exists a legal tiling τ. Conditions 1, 2 and 3 ensure that the REM in (3) defines $\{\langle p_2, q_2\rangle\}$. Conversely, suppose $\{\langle p_2, q_2\rangle\}$ is definable. There exists a defining REM e and a data path w in $\mathcal{L}(e)$ connecting p_2 to q_2. If w does not encode a legal tiling, then condition 4 above implies that there is a data path w' automorphic to w (and hence in $\mathcal{L}(e)$) connecting p_1 to q_1, contradicting the hypothesis that e defines $\{\langle p_2, q_2\rangle\}$. Hence, w encodes a legal tiling. The data graph can be constructed in polynomial time (details follow) and hence the RDPQ_{mem}-definability problem in data graphs is EXPSPACE-hard. At a high level, the strategy of this proof is similar to that of [17, Theorem 3.7] in the sense that a small gadget differentiates between the set of all tilings and the set of illegal tilings. However, [17, Theorem 3.7] can not be used here directly, since that is about containment of one query in another while we are concerned about the definability of a relation in a given data graph. There is also a subtle difference between the proof strategies which will be highlighted in the details that follow.

We now give the details of the data graph. Nodes are denoted by circles, with data values written outside. The data values of some nodes are skipped when they are not important. The data values $d_n, e_n, \ldots, d_1, e_1$ are all distinct. The portion of the data graph containing p_2 and q_2 is as follows.

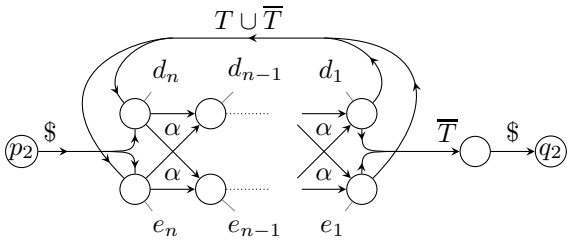

For every data path connecting p_2 to q_2 that is not the encoding of a legal tiling, we now add automorphic data paths connecting p_1 to q_1 through gadgets. This will ensure that the data graph satisfies condition 4 above. We also ensure that every path we add satisfies conditions 1 and 3.

- *In a data path w connecting p_2 to q_2, the sequence of n data values preceding $\tau(0,1)$ does not represent* 1. This could be due to any one (or more) of the n bits being wrong; following is the gadget for checking that the k^{th} bit is 1 (at node q) instead of 0. There are n such gadgets, one for each bit.

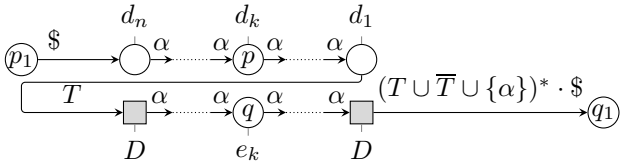

The gadget simply checks that the k^{th} data value preceding $\tau(0,1)$ in the data path w (e_k at node q) is unequal to the k^{th} data value preceding $\tau(0,0)$ (d_k at node p). This will ensure that the k^{th} bit preceding $\tau(0,1)$ is 1 instead of 0. If the k^{th} data value preceding $\tau(0,0)$ is e_k, the gadget can still imitate w modulo an automorphism that interchanges d_k and e_k. In the above diagram, every gray box marked D is actually a gadget with $2n$ nodes, with each node having a distinct data value from the set $\{d_n, e_n, \ldots, d_1, e_1\}$. For every edge coming in to a gray box, there is an edge coming in to each of the $2n$ nodes. For every edge coming out of a gray box, there is an edge coming out of each of the $2n$ nodes. The edge coming in to the node q_1 is labeled by the REM $(T \cup \overline{T} \cup \{\alpha\})^* \cdot \$$, whose language consists of all the data paths having exactly one occurrence of the letter \$, which occurs at the end. A gadget admitting exactly this set of data paths can be easily designed using polynomially many nodes; the gadget is not shown in the diagram since it is easier to understand the expression.

- *Some sequence of n conditions does not encode the successor of the preceding n conditions.* The following gadget checks that the k^{th} bit flips from 1 (in node 2) to 0 (in node 4) but the $(k+1)^{\text{th}}$ bit stays at 0 (in nodes 1 and 3). There are $\mathcal{O}(n)$ such gadgets for checking all such errors.

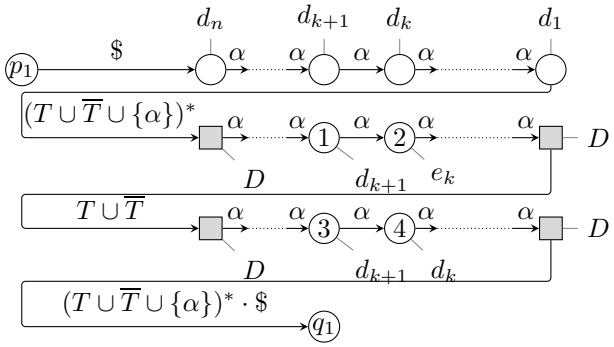

- *The sequence of n conditions before a letter in $\overline{T}$ does not represent* $2^n - 1$. This can be due to any one (or more) of the bits being 0 instead of 1. The following gadget checks that the k^{th} bit is 0 (at node q). There are n

such gadgets, one for each bit.

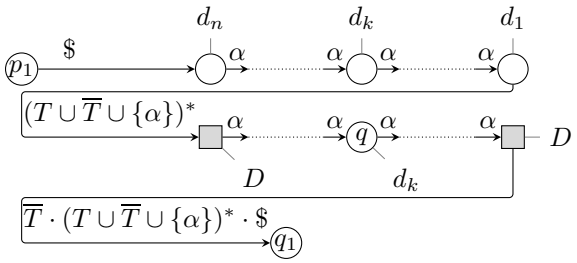

- *The n conditions before a letter in T represent $2^n - 1$.*

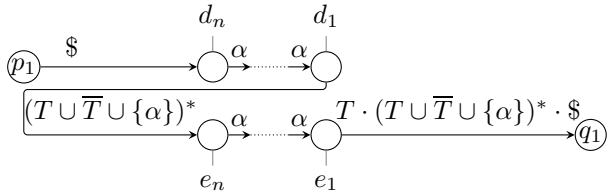

- *The tiling does not begin with the tile type t_i.*

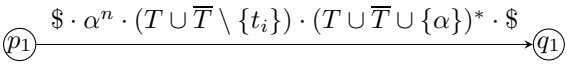

- *The tiling does not end with the tile type t_f.*

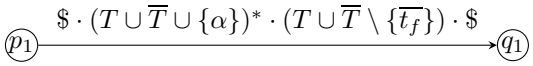

- *Two adjacent tiles in the same row are not horizontally compatible.* The following gadget checks that the tile type t_2 is adjacent to the horizontally incompatible tile type t_1 in the same row. There is one such gadget for every pair (t_1, t_2) of horizontally incompatible tile types.

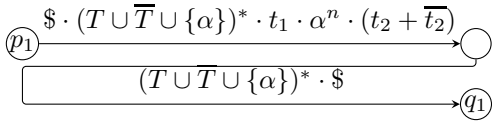

- *Two adjacent tiles in the same column are not vertically compatible.* The gadget below checks that the tile type t_2 (seen just after node 2) is adjacent to the vertically incompatible tile type t_1(seen just after node 1) in the last column. The tiles are matched from the same column, since the data values seen just before $\overline{t_1}$ are same as the data values seen just before $\overline{t_2}$. The tiles are matched from adjacent rows, since $\overline{t_1}$ is the only letter from $\overline{T}$ allowed between the nodes 1 and 2. There is one such gadget for every pair (t_1, t_2) of vertically incompatible tile types.

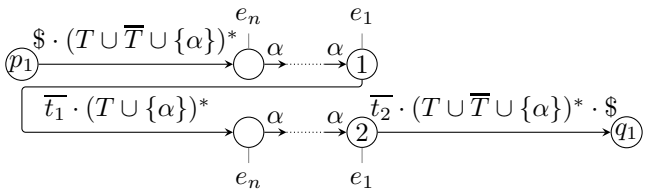

We now highlight a subtle difference between this proof and that of [17, Theorem 3.7]. There, to check that the distance between two positions is exactly 2^n, a gadget checks n bits of the first position against n bits of the second position. There the gadget is built using REM and hence it can check the bits individually for equality. Here, we need to build a similar gadget using data graphs. If we did this check by matching each bit explicitly, the gadget would be exponentially larger. We avoid it by observing that we need not admit the exact data path encoding an illegal tiling, but only an *automorphic copy*. Thus, if a data path has the data value d_1 just before $\overline{t_1}$ and $\overline{t_2}$, the above gadget will still catch it through an automorphism that interchanges d_1 and e_1.

The gadget below checks that the tile type t_2 is adjacent to the vertically incompatible tile type t_1 in a column other than the last one. There is one such gadget for every pair (t_1, t_2) of vertically incompatible tile types.

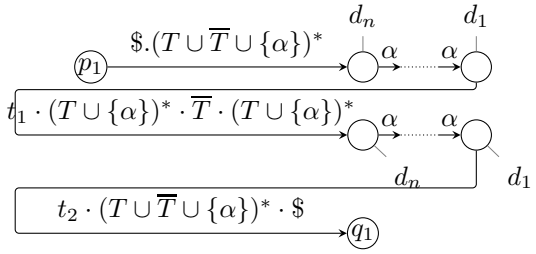

□

4 Queries using Regular expressions with equality

We study the RDPQ$_=$-definability problem in this section. These are queries using REE. As seen before, REE have the additional $e_=$ and $e_{\neq}$ constructs on top of standard regular expressions. In Example 12, we have seen that they are less powerful than REM in defining relations. However, due to $e_=$ and $e_{\neq}$, they can define more relations than RPQs.

In Section 3, we have shown that RDPQ$_{\text{mem}}$-definability is EXPSPACE-complete. RPQ-definability is known to be PSPACE-complete [3]. We will now prove that RDPQ$_=$-definability is PSPACE-complete as well.

The main idea is the following observation. Suppose we have an expression $e_1 \cdot e_2$. If e_1 and e_2 are REMs, it is possible that there is a register getting bound in e_1 by $\downarrow r$, and used in a condition in e_2. So one cannot reason about $e_1 \cdot e_2$ by independently reasoning about e_1 and e_2. Such a situation does not arise with REE. The relation defined by $e_1 \cdot e_2$ can in fact be obtained as a composition of the relations defined by e_1 and e_2. This makes it possible to solve the problem in PSPACE.

For this section, fix a data graph $G = (V, E, \rho)$ over edge alphabet Σ. Let $\mathsf{Bin} = V \times V$ be the set of binary relations over V. Our first goal would be to define operators over the set Bin and to generate relations in a hierarchical manner by making use of these operators in a certain way.

Definition 26. *Given two relations $S_1, S_2 \in \mathsf{Bin}$, we define the following operators:*

$$\begin{aligned} S_1 + S_2 &= \{\langle u, v\rangle \mid \langle u, v\rangle \in S_1 \text{ or } \langle u, v\rangle \in S_2\} \\ S_1 \circ S_2 &= \{\langle u, v\rangle \mid \exists z : \langle u, z\rangle \in S_1 \text{ and } \langle z, v\rangle \in S_2\} \\ S^{=} &= \{\langle u, v\rangle \in S \mid \rho(u) = \rho(v)\} \\ S^{\neq} &= \{\langle u, v\rangle \in S \mid \rho(u) \neq \rho(v)\} \end{aligned}$$

The $+$ and $\circ$ are the *union* and *composition* operators. We call $S^{=}$ and $S^{\neq}$ as the *$=$-restriction* and the *$\neq$-restriction* respectively. From the definitions, it is easy to see that $+$ is commutative and associative. Furthermore, the operator $\circ$ is associative and distributes over $+$.

For an REE e, let us write S_e for the binary relation defined by it: $S_e = \{\langle u, v\rangle \in V \times V \mid \exists w \in \mathcal{L}(e) \text{ s.t. } u \xrightarrow{w} v\}$.

Definition 27. *Consider the set* Bin *of binary relations. We will define a sequence* $L_0, L_1, \ldots$ *of subsets of* Bin*, called* levels*, as follows:*

$$\begin{aligned} B_0 &= \{S_\varepsilon\} \cup \{S_a \mid a \in \Sigma\} \\ L_0 &= \textit{closure of } B_0 \textit{ under } + \textit{ and } \circ \\ \textit{for } i \geq 1,\ B_i &= \{S^{=} \mid S \in L_{i-1}\} \cup \{S^{\neq} \mid S \in L_{i-1}\} \cup L_{i-1} \\ L_i &= \textit{closure of } B_i \textit{ under } + \textit{ and } \circ \end{aligned}$$

Intuitively, the set B_0 consists of those relations that are defined using the atomic expressions ε and a. The set L_0 closes these relations under union and composition. The base sets B_1 of the next level are formed by adding to L_0 the $=$ and $\neq$-restrictions of relations in L_0. These are now closed under union and composition to get the set L_1. This process continues. Of course, this cannot continue beyond 2^{n^2} steps, which is the total number of relations in Bin. The next lemma further restricts it to n^2 steps.

Lemma 28. *For all* $j \geq n^2$, $L_j = L_{n^2}$.

Proof. To every newly added relation in L_i, we will associate a relation in L_{i-1} having at least one extra pair of nodes. As the new relations that are added become strictly smaller each time we go up a level, no new relations can be added beyond L_{n^2}.

A newly added relation S in L_i is either in B_i, or formed by union of compositions of relations from B_i:

$$\begin{aligned} S &= T_1 + T_2 + \cdots + T_m && \text{where } m \leq n^2 \\ \text{and each } T_j &= (R_1 \circ R_2 \circ \ldots \circ R_p) && \text{s.t. } p \leq 2^{n^2}, \text{ and} \\ & && R_k \in B_i \text{ for } k \leq p. \end{aligned} \tag{4}$$

The bounds on m and p above follow from the fact that each relation can have at most n^2 pairs. For the relations $S^{=}$ or $S^{\neq}$ added in B_i let us associate the set $S \in L_{i-1}$. They are added only if they are strict subsets of S. This means there exists a pair $\langle u, v\rangle \in S$ that does not belong to $S^{=}$ and another pair $\langle u', v'\rangle \in S$ that does not belong to $S^{\neq}$. Hence the cardinalities of $S^{=}$ and $S^{\neq}$ are strictly lesser than S.

Let T be a relation obtained by union of compositions of relations in B_i. The relation T is new only if some of the relations in the underlying composition according to (4) are the new basic sets $S^{=}$ and $S^{\neq}$. Consider the relation T' where each of these $S^{=}$ and $S^{\neq}$ is replaced by the corresponding relation S. Clearly $T \subseteq T'$. Moreover T is added only if it is different from T'. This shows that T' has at least one pair more than T. □

The motive behind defining these operations and the hierarchy of relations is that this procedure resembles the way REEs are constructed from its grammar.

Lemma 29. *For every two REE* e *and* f*, we have:* $S_e + S_f = S_{e+f}$, $S_e \circ S_f = S_{ef}$, $S_e^{=} = S_{e_=}$ *and* $S_e^{\neq} = S_{e_\neq}$.

The proof of the above lemma is quite straightforward from the definitions. However, the lemma is significant because it allows to reason about S_{ef} by independently reasoning about S_e and S_f. As mentioned before, this property is not true for REMs. The above lemma can be used to show the important property that all REE-definable relations can be generated by this hierarchical construction that repeatedly applies the $=$ and $\neq$ restrictions and closes under $+$ and $\circ$.

Lemma 30. *A relation is* $\mathrm{RDPQ}_{=}$*-definable iff it belongs to level* L_{n^2}.

Proof. We prove the left-to-right direction by an induction on the structure of REE. Relations S_ε and S_a definable by the basic REE ε and a already belong to L_0, and hence belong to L_{n^2}. For REE e and f, let S_e and S_f belong to L_{n^2}.

By Lemma 29, we have $S_{e+f} = S_e + S_f$ and $S_{ef} = S_e \circ S_f$. As L_{n^2} is closed under $+$ and $\circ$, the relations S_{e+f} and S_{ef} belong to L_{n^2} as well. The relation $S_{e_=}$ equals $S_e^{=}$. Since S_e belongs to L_{n^2}, the relation $S_e^{=}$ would be

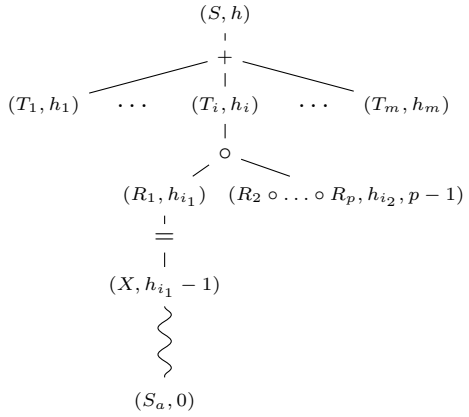

Figure 2: A part of execution of the algorithm for RDPQ$_=$ definability

present in B_{n^2+1} by definition and hence in the level L_{n^2+1}. But by Lemma 28, this means that $S_e^=$ belongs to L_{n^2} as well. Similar argument holds for $S_e^{\neq}$.

The right-to-left direction can be proved by an easy induction on the level number, once again using Lemma 29. □

Let us define the *height* of a relation S to be the least i such that $S \in L_i$. The fact that the height of an RDPQ$_=$ definable relation is polynomially bounded can be used to give a PSPACE upper bound.

Lemma 31. RDPQ$_=$ *definability problem is in* PSPACE.

Proof. The inputs are a data graph G and a binary relation S. We will describe a non-deterministic algorithm that decides in polynomial space if S is RDPQ$_=$-definable. Due to Lemma 30, checking if S is RDPQ$_=$ definable is equivalent to checking if S belongs to L_{n^2}. We will now explain how the algorithm can check the membership of S in L_{n^2}.

Every relation in L_{n^2} has appeared due to a sequence of computations starting from the basic relations S_ε and S_a for each $a \in \Sigma$. Thanks to (4), there is a specific structure to this computation that allows to look at it as a "computation tree". Nodes in this tree are relations in L_{n^2} and the children of a node are different relations in the same or smaller level that are used to construct the parent relation through the operations $+$, $\circ$, or the $=$ and $\neq$ restrictions. The leaves are the basic sets S_ε and S_a.

A naive approach would be to guess this entire computation tree and check two things: are leaves of the form S_a or S_ε, and is each node either the union, composition or one of the (in)equality restrictions of its children. These checks can be done in polynomial time. However, we cannot hope to maintain the entire tree in PSPACE as compositions can have exponentially many children (4).

Instead of guessing the entire tree, the algorithm guesses, in a certain way, a path in the tree along with the children of each node in this path: if the node is a union of $T_1, \ldots, T_m$, all these relations are guessed as children (there are only polynomially many); if a node is a composition of $R_1, \ldots, R_p$, two children are guessed - the relation R_1 and the composed relation $R_2 \circ \cdots \circ R_p$. The former is a proper child in the computation tree, and the latter is a different relation which needs to be further decomposed to get the actual children in the computation tree. The algorithm also maintains the number the decompositions left: so the child would be $(R_2 \circ \cdots \circ R_p, p-1)$. Storing $p-1$ needs only polynomially many bits. This is the main idea. Additionally, the algorithm maintains the height of each node (see Figure 2). Each time an $=$ or $\neq$ restriction happens, the height reduces by one. At the leaf level, it is checked if the relation is one of S_ε or S_a. Nodes whose heights have been certified can be removed: that is, leaf nodes can be removed after the basic check, and non-leaf nodes are removed once all its children are checked.

Hence at any point of time, the algorithm maintains some structure like Figure 2. Since the height is polynomial, this can be done in PSPACE. □

The PSPACE-hardness of RPQ-definability [3] can be easily extended to give the same lower bound for RDPQ$_=$-definability.

Theorem 32. RDPQ$_=$*-definability is* PSPACE*-complete.*

Proof. The PSPACE upper bound comes from Lemma 31. For the lower bound, consider the RPQ-definability problem: given a graph H and a relation T on H, is T definable by a regular expression? This problem is known to be PSPACE-complete [3]. We will show that RPQ-definability can be reduced to RDPQ$_=$ definability.

Consider the data graph H' obtained from H by attaching the same data value to all the nodes. Consider the RDPQ$_=$ definability of the same set T. If T is RPQ-definable on H, then clearly T would be RDPQ$_=$ definable on H' as well using the same expression. Suppose T is RDPQ$_=$ definable on H' using the expression e. Without loss of generality, assume that T is non-empty. We claim that the REE defining T on H' will not have sub-expressions of the form $f_{\neq}$. This is because $f_{\neq}$ defines the empty relation and hence can be eliminated from the defining REE. Secondly observe that for the graph H', no matter which REE f we choose, $S_f = S_{f_=}$. Hence all sub-expressions of the form $f_=$ in e can be modified to f and still the defined set remains the same. This way, we have obtained a regular expression that defines H'. The same regular expression would define T on H as well.

□

5 Union of Conjunctive Queries

In this section, we study the complexity of the definability problem for data graphs using UCRDPQs. The notion of homomorphisms has been used in relational databases to characterize relations definable by union of conjunctive queries. We will now adapt it to data graphs.

Definition 33. *Let $G = (V, E, \rho)$ be a data graph and $h : V \to V$ be a mapping. We call h a* data graph homomorphism *if it satisfies the following two conditions.*

1. *(Single step compatibility) For every $p, q \in V$ and $a \in \Sigma$, $p \xrightarrow{a} q$ implies $h(p) \xrightarrow{a} h(q)$.*
2. *(Data compatibility of reachable nodes) For every $p, q \in V$, if q is reachable from p, then $\rho(p) = \rho(q) \Leftrightarrow \rho(h(p)) = \rho(h(q))$.*

Intuitively, a data graph homomorphism h ensures that if there is an edge labeled a from p to q, there is also an edge labeled a from $h(p)$ to $h(q)$, thus preserving the relations induced by the letters in the finite alphabet. In addition, suppose there is a path from p to q. Then the data values at p and q are same if, and only if, the data values at $h(p)$ and $h(q)$ are same. This preserves the relations induced by (in)equality of data values. The following result characterizes UCRDPQ-definable sets in terms of data graph homomorphisms.

Lemma 34. *Let $G = (V, E, \rho)$ be a data graph and S be a relation of any arity. Then the following are equivalent.*

1. *The set S is UCRDPQ-definable.*
2. *For every data graph homomorphism h and every tuple $\overline{p} \in S$, $h(\overline{p})$ also belongs to S.*

Proof. $(1 \Rightarrow 2)$. Suppose $S = Q_1(G) \cup \cdots \cup Q_k(G)$, where $Q_1, \ldots, Q_k$ are CRDPQs. Suppose Q_j is of the form (1)(in Definition 13) and $\overline{p} \in Q_j(G)$. Let h be a data graph homomorphism. We will prove that $h(\overline{p}) \in Q_j(G)$. From the semantics of CRDPQs, we infer that there is a valuation $\mu : \cup_{1 \leq i \leq m}\{x_i, y_i\} \to V$ such that $(G, \mu) \models Q_j$ and $\overline{p} = \mu(\overline{z})$. Let $h \circ \mu$ be the valuation such that $h \circ \mu(x) = h(\mu(x))$ for all x. It is enough to prove that $(G, h \circ \mu) \models Q_j$ — in that case, $h(\overline{p}) = h \circ \mu(\overline{z}) \in Q_j(G)$. Now let us prove that $(G, h \circ \mu) \models Q_j$. Since Q_j is of the form (1) and $(G, \mu) \models Q_j$, we infer that for every $i = 1, \ldots, m$, there is a data path $w_i \in \mathcal{L}(e_i)$ from

$\mu(x_i)$ to $\mu(y_i)$. It is enough to prove that there is a data path w_i' automorphic to w_i from $h \circ \mu(x_i)$ to $h \circ \mu(y_i)$ for every $i = 1, \ldots, m$ — in that case $(G, h \circ \mu) \models Q_j$. So suppose w_i is the data path associated with the path $\mu(x_i) \xrightarrow{a_1} p_2 \xrightarrow{a_2} \cdots \xrightarrow{a_{l-1}} p_l \xrightarrow{a_l} \mu(y_i)$, where $p_2, \ldots, p_l$ are the intermediate nodes in the path from $\mu(x_i) = p_1$ to $\mu(y_i) = p_{l+1}$. Since h is a data graph homomorphism, we infer from the single step compatibility property that G has a path $h \circ \mu(x_i) \xrightarrow{a_1} h(p_2) \xrightarrow{a_2} \cdots \xrightarrow{a_{l-1}} h(p_l) \xrightarrow{a_l} h \circ \mu(y_i)$, where $h(p_2), \ldots, h(p_l)$ are intermediate nodes in a path from $h \circ \mu(x_i)$ to $h \circ \mu(y_i)$. We claim that the data path w_i' associated with this path is automorphic to w_i. If not, there would be positions j_1 and j_2 such that $\rho(p_{j_1}) = \rho(p_{j_2})$ but $\rho(h(p_{j_1})) \neq \rho(h(p_{j_2}))$ (or vice-versa), violating the data compatibility property of the data graph homomorphism h. Thus, w_i' is a data path automorphic to w_i from $h \circ \mu(x_i)$ to $h \circ \mu(y_i)$. This concludes the proof that $h(\overline{p}) \in S$.

$(2 \Rightarrow 1)$ Suppose $V = \{p_1, \ldots, p_n\}$. Let $\overline{x} = \langle x_1, \ldots, x_n \rangle$. Let $\phi_G(\overline{x})$ be defined as follows.

$$\phi_G(\overline{x}) = \bigwedge_{(p_i, a, p_j) \in E} x_i \xrightarrow{a} x_j \quad \wedge \bigwedge_{(p_i, p_j) \in (\Sigma^+)_=(G)} x_i \xrightarrow{(\Sigma^+)_=} x_j$$
$$\wedge \bigwedge_{(p_i, p_j) \in (\Sigma^+)_{\neq}(G)} x_i \xrightarrow{(\Sigma^+)_{\neq}} x_j$$

In the above definition, $(\Sigma^+)_=$ is an REE; $(\Sigma^+)_=(G)$ is the set of pairs of nodes (p_i, p_j) such that p_j is reachable from p_i and both nodes have the same data value. Similarly, $(\Sigma^+)_{\neq}(G)$ is the set of pairs of nodes (p_i, p_j) such that p_j is reachable from p_i and the two nodes have different data values. The valuation that assigns p_i to x_i for every $i = 1, \ldots, n$ satisfies all the conditions in $\phi_G(\overline{x})$. If a valuation μ for $\overline{x}$ satisfies all the conditions in $\phi_G(\overline{x})$, then the mapping $h_\mu : V \to V$ such that $h_\mu(p_i) = \mu(x_i)$ is a data graph homomorphism. Let S' be the set of tuples defined by the UCRDPQ $\{Ans(\langle x_{i_1}, \ldots, x_{i_r} \rangle) := \phi_G(\overline{x}) \mid \langle p_{i_1}, \ldots, p_{i_r} \rangle \in S\}$. We claim that $S' = S$, which will prove that S is UCRDPQ-definable.

$(S \subseteq S')$: For every $\langle p_{i_1}, \ldots, p_{i_r} \rangle \in S$, we have that $\langle p_{i_1}, \ldots, p_{i_r} \rangle \in Ans(\langle x_{i_1}, \ldots, x_{i_r} \rangle) := \phi_G(\overline{x})(G) \subseteq S'$.

$(S' \subseteq S)$: Suppose $\langle p_{j_1}, \ldots, p_{j_r} \rangle \in S'$. Then there is some $\langle p_{i_1}, \ldots, p_{i_r} \rangle \in S$ and a valuation μ for $\overline{x}$ such that μ satisfies all the conditions in $\phi_G(\overline{x})$ and $\langle p_{j_1}, \ldots, p_{j_r} \rangle = \langle\, \mu(x_{i_1}), \ldots, \mu(x_{i_r}) \rangle$. The mapping $h_\mu : V \to V$ such that $h_\mu(p_i) = \mu(x_i)$ is a data graph homomorphism. Now we have $\langle p_{j_1}, \ldots, p_{j_r} \rangle = \langle \mu(x_{i_1}), \ldots, \mu(x_{i_r}) \rangle$ and in addition $\langle \mu(x_{i_1}), \ldots, \mu(x_{i_r}) \rangle = \langle h_\mu(p_{i_1}), \ldots, h_\mu(p_{i_r}) \rangle$. Hence, we have $\langle p_{j_1}, \ldots, p_{j_r} \rangle = h_\mu(\langle p_{i_1}, \ldots, p_{i_r} \rangle) \in S$; the last inclusion follows from condition 2 of the lemma, as $\langle p_{i_1}, \ldots, p_{i_r} \rangle \in S$. □

Readers familiar with Global as View (GAV) schema mappings for relational databases may note some similarities with the above result. For a data graph G, consider the relational database D_G over the domain V consisting of all the binary relations that are definable by RDPQs. For a set of tuples S over V, let D_S be the relational database consisting of the single relation S. Then S is UCRDPQ-definable on G iff some GAV schema mapping fits the source database D_G and the target database D_S. A characterization using homomorphisms similar to the one in Lemma 34 is given for GAV schema mappings in [9, 2]. Here, we extend the notion of homomorphisms to include data value compatibility. A coNP-completeness result for a subclass of GAV schema mappings is given in [9]. We give a similar result for UCRDPQ-definability below. However, there is no obvious way of directly using the upper bound in [9] here, since the relational database D_G may be exponentially larger than G. Preservation under homomorphism is a fundamental concept, which appears in other contexts as well, for example querying databases with incomplete information [12].

Theorem 35. *UCRDPQ-definability is* coNP*-complete.*

Proof. We first prove the coNP upper bound. Given a data graph $G = (V, E, \rho)$ and a set of tuples S, we can guess a mapping $h : V \to V$, verify that it is a data graph homomorphism and that there is some tuple $\overline{p} \in S$ such that $h(\overline{p}) \notin S$. If S is not UCRDPQ-definable, then Lemma 34 ensures that at least one of the guesses will succeed. On the other hand, if S is UCRDPQ-definable, then none of the guesses will succeed.

For the coNP lower bound, we reduce the unsatisfiability problem for Boolean 3-CNF formulas to the UCRDPQ-definability problem. This part of the proof is an adaptation of a similar proof from [9] about relational databases

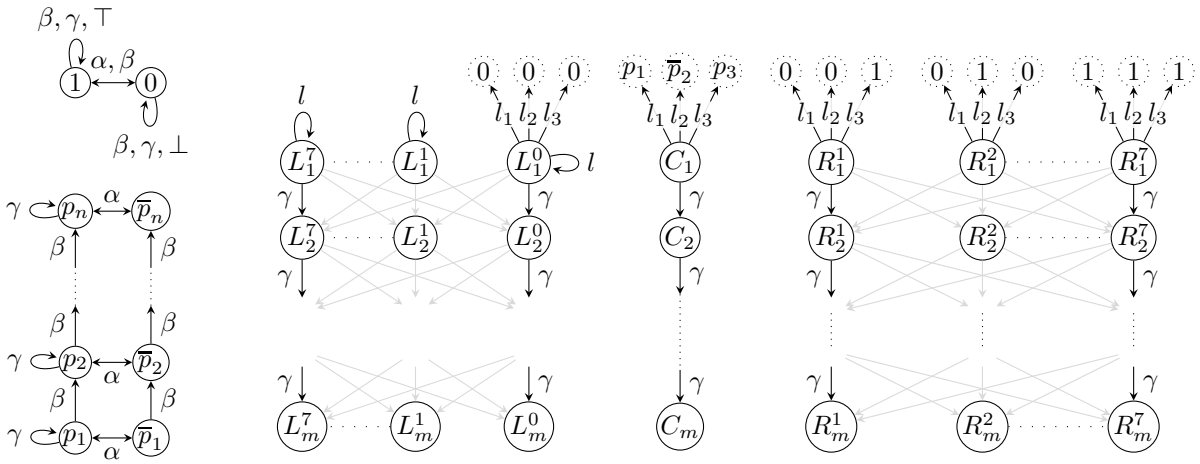

Figure 3: Data graph for the lower bound of UCRDPQ-definability

(which can have ternary relations) to data graphs (which can only have binary relations). Given a 3-CNF formula F consisting of clauses $C_1, \ldots, C_m$ over the variables $p_1, \ldots, p_n$, we map it to the data graph shown in Fig. 3. All nodes have the same data value, which is not shown explicitly.

There is only one node $\textcircled{1}$, but the same node is drawn (dotted) at many places in the diagram to avoid the clutter of too many edges. Same comment applies to all the nodes drawn dotted. There is an edge labeled γ from all nodes R_i^j to R_{i+1}^k, but most of the edges are grayed out and the label γ is not shown to reduce clutter. Same applies to edges from L_i^j to L_{i+1}^k. In the diagram, the clause C_1 is assumed to be $(p_1 \vee \neg p_2 \vee p_3)$. From every clause node C_i, there are edges labeled l_1, l_2 and l_3 to the nodes corresponding to the literals occurring in C_i. Only the edges from C_1 are shown and others are not shown. From every node R_i^j and L_i^j, there is an edge labeled l_k to either 0 or 1, depending on the k^{th} most significant bit of the binary representation of j. Only the edges from L_1^0, R_1^1, R_1^2 and R_1^7 are shown. Others are not shown. We claim that the given Boolean 3-CNF formula F is not satisfiable iff the set of tuples of nodes $S = \{\langle C_1\rangle, \ldots, \langle C_m\rangle\} \cup \{\langle L_i^j\rangle \mid 1 \le i \le m, 0 \le j \le 7\}$ is UCRDPQ-definable.

Suppose there is an assignment $sa : \{p_1, \ldots, p_n\} \to \{0, 1\}$ satisfying F. Consider the graph mapping h that maps the node p_i to the node $sa(p_i)$ and $\overline{p}_i$ to $1 - sa(p_i)$. For every $i = 1, \ldots, m$, h maps the node C_i to the node R_i^j; here j is the number whose binary representation is the one formed by the three literals of the clause C_i according to the satisfying assignment sa. All other nodes are mapped to themselves by h. This mapping h is a data graph homomorphism and $h(\langle C_1\rangle) = \langle R_1^j\rangle$ for some j. Since $\langle C_1\rangle \in S$ and $\langle R_1^j\rangle \notin S$, we infer from Lemma 34 that S is not UCRDPQ-definable.

Conversely, suppose F is not satisfiable. Let h be any data graph homomorphism. We will prove that $h(\langle p\rangle) \in S$ for every tuple $\langle p\rangle \in S$. Since the only node with a self edge labeled $\top$ (resp. $\bot$) is 1 (resp. 0), 1 (resp. 0) is mapped to itself by h. Due to the self edges labeled l and the edges labeled l_1, l_2 and l_3, h maps L_1^j to itself for every $j = 0, \ldots, 7$. The edges labeled γ, l_1, l_2, l_3 then force h to map L_i^j to itself for every i, j. It remains to prove that $h(\langle C_i\rangle) \in S$ for every $i = 1, \ldots, m$. Due to the edges labeled β and the self edges labeled γ, h maps p_1 to either itself or to 1 or to 0. If h maps p_1 to itself, then the edges labeled α and β force h to map p_i to itself (and $\overline{p}_i$ to itself) for every $i = 1, \ldots, n$. The edges labeled l_1, l_2 and l_3 then force h to map C_i to itself for every $i = 1, \ldots, m$. On the other hand, if h maps p_1 to 1 or 0, the edges labeled α force h to map $\overline{p}_1$ to $1 - h(p_1)$. The edges labeled α and β then force h to map p_i to 1 or 0 and $\overline{p}_i$ to $1 - h(p_i)$ for every $i = 1, \ldots, n$. The homomorphism h thus determines a truth assignment for $p_1, \ldots, p_n$. For every $i = 1, \ldots, m$, the edges labeled l_1, l_2 and l_3 force h to map C_i to either L_i^j or R_i^j; here j is the number whose binary representation is the one formed by the three literals of the clause C_i according to the truth assignment determined by h. If h maps C_i to R_i^* (R_i^* could be any one of $R_i^1, \ldots, R_i^7$) for some $i = 1, \ldots, m$, then the edges labeled γ force h to map C_i to R_i^* for every $i = 1, \ldots, m$. This

implies that the truth assignment determined by h assigns at least one literal to *true* in every clause, contradicting the hypothesis that F is not satisfiable. Hence h maps C_i to L_i^* for every $i = 1, \ldots, m$. Since this holds for every data graph homomorphism, we conclude that $h(\langle p \rangle) \in S$ for every data graph homomorphism h and every tuple $\langle p \rangle \in S$. Hence, we can conclude from Lemma 34 that S is UCRDPQ-definable. □

6 Discussion

A natural question to ask is how to synthesize a query that defines a given relation. In principle, the decision procedures in the paper can be converted into a procedure to synthesize a defining query. However such queries would not have an interesting structure. For instance, in the REM and REE cases, the synthesized queries do not make use of the star operator. Moreover, the lower bound for the decision problem implies that the worst case size of the defining queries will be doubly exponential for REMs, and exponential for REEs. In the UCRDPQ case, the defining query described in Lemma 34 essentially constructs the whole data graph using conditions and then picks out the required tuples. This does not capture the essence of conjunctive queries, which is to identify patterns that are much smaller than the graphs themselves.

A possible future direction would be to find a notion of "good" queries and reformulate the definability problem to ask for the existence of "good" defining queries.

In some application domains, data graphs may have a special structure (such as not too many cycles). An orthogonal direction would be to study the definability problem for such data graphs.

References

[1] B. Alexe, B. T. Cate, P. G. Kolaitis, and W-C Tan. Characterizing schema mappings via data examples. *ACM Trans. Database Syst.*, 36(4):23:1–23:48, 2011.

[2] B. Alexe, B. T. Cate, P. G. Kolaitis, and W-C Tan. Designing and refining schema mappings via data examples. In *SIGMOD*, pages 133–144, 2011.

[3] T. Antonopoulos, F. Neven, and F. Servais. Definability problems for graph query languages. In *ICDT*, pages 141–152, 2013.

[4] P. Barceló. Querying graph databases. In *PODS*, pages 175–188, 2013.

[5] P. Barceló, J. Pérez, and J. Reutter. Schema mappings and data exchange for graph databases. In *ICDT*, pages 189–200, 2013.

[6] P. Barceló, J. Reutter, and L. Libkin. Parameterized regular expressions and their languages. *Theor. Comput. Sci.*, 474:21–45, 2013.

[7] D. Calvanese, G. De Giacomo, M. Lenzerini, and M. Y. Vardi. Simplifying schema mappings. In *ICDT*, pages 114–125, 2011.

[8] B. T. Cate, V. Dalmau, and P. G. Kolaitis. Learning schema mappings. *ACM Trans. Database Syst.*, 38(4):28:1–28:31, 2013.

[9] B. T. Cate, P. G. Kolaitis, and W-C Tan. Database constraints and homomorphism dualities. In *CP*, pages 475–490, 2010.

[10] A. Das Sarma, A. Parameswaran, H. Garcia-Molina, and J. Widom. Synthesizing view definitions from data. In *ICDT*, pages 89–103, 2010.

[11] G.H.L. Fletcher, M. Gyssens, J. Paredaens, and D. V. Gucht. On the expressive power of the relational algebra on finite sets of relation pairs. *IEEE Trans. Knowledge and Data Engg.*, 21(6):939–942, 2009.

[12] A. Gheerbrant, L. Libkin, and C. Sirangelo. When is naive evaluation possible? In *PODS*, pages 75–86, 2013.

[13] G. Gottlob, C. Koch, R. Baumgartner, M. Herzog, and S. Flesca. The Lixto data extraction project: Back and forth between theory and practice. In *PODS*, pages 1–12, 2004.

[14] G. Gottlob and P. Senellart. Schema mapping discovery from data instances. *J. ACM*, 57(2):6:1–6:37, 2010.

[15] C. Gutierrez, C. Hurtado, and A. Mendelzon. Foundations of semantic web databases. *JCSS*, 77(3):520–541, 2011.

[16] M. Kaminski and N. Francez. Finite-memory automata. *Theor. Comp. Sc.*, 134(2):329 – 363, 1994.

[17] E.V. Kostylev, J.L. Reutter, and D. Vrgoč. Containment of data graph queries. In *ICDT*, pages 131–142, 2014.

[18] U. Leser. A query language for biological networks. *Bioinformatics*, 21(suppl 2):ii33–ii39, 2005.

[19] L. Libkin and D. Vrgoč. Regular expressions for data words. In *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 7180 of *LNCS*, pages 274–288. 2012.

[20] L. Libkin and D. Vrgoč. Regular path queries on graphs with data. In *ICDT*, pages 74–85, 2012.

[21] F. Neven, T. Schwentick, and V. Vianu. Finite state machines for strings over infinite alphabets. *ACM TOCL*, 5(3):403–435, 2004.

[22] J. Pérez, M. Arenas, and C. Gutierrez. Semantics and complexity of SPARQL. *ACM Trans. Database Syst.*, 34(3), 2009.

[23] R. Ronen and O. Shmueli. Soql: A language for querying and creating data in social networks. In *ICDE*, pages 1595–1602, 2009.

[24] H. Sakamoto and D. Ikeda. Intractability of decision problems for finite-memory automata. *Theor. Comp. Sc.*, 231(2):297 – 308, 2000.

[25] P. van Emde Boas. The convenience of tilings. In *Complexity, Logic and Recursion Theory*, pages 331–363. Marcel Dekker Inc., 1997.